arXiv:1607.06334v1 [math.GR] 21 Jul 2016

CLASSIFYING VIRTUALLY SPECIAL TUBULAR GROUPS

DANIEL J. WOODHOUSE

Abstract. A group is tubular if it acts on a tree with $\mathbb{Z}^2$ vertex stabilizers and $\mathbb{Z}$ edge stabilizers. We prove that a tubular group is virtually special if and only if it acts freely on a locally finite CAT(0) cube complex. Furthermore, we prove that if a tubular group acts freely on a finite dimensional CAT(0) cube complex, then it virtually acts freely on a three dimensional CAT(0) cube complex.

1. Introduction

A *tubular group* G splits as a graph of groups with $\mathbb{Z}^2$ vertex groups and $\mathbb{Z}$ edge groups. Equivalently, G is the fundamental group of a graph of spaces, denoted by X, with each vertex space homeomorphic to a torus and each edge space homeomorphic to $S^1 \times [-1, 1]$. The graph of spaces X is a *tubular space*. In this paper all tubular groups will be finitely generated and therefore have compact tubular spaces.

Tubular groups have been studied from various persectives: Brady and Bridson provided tubular groups with isoperimetric function n^α for all α in a dense subset of $[2, \infty)$ in [1]. Cashen determined when two tubular groups are quasi-isometric [5]. Wise determined whether or not a tubular group acts freely on a CAT(0) cube complex [11], and classified which tubular groups are cocompactly cubulated. The author determined a criterion for finite dimensional cubulations [13]. Button has proven that all free-by-cyclic groups that are also tubular groups act freely on finite dimensional cube complexes [3]. The main theorem of this paper is:

Theorem 5.10. *A tubular group G acts freely on a locally finite CAT(0) cube complex if and only if G is virtually special.*

Haglund and Wise introduced special cube complexes in [7]. The main consequence of a group being special is that it embeds in a right angled Artin group (see [9] or [10] for a full outline of Wise's program).

1.1. **Structure of the Paper.** In [11] Wise obtained free actions of tubular groups on CAT(0) cube complexes by first finding *equitable sets* that allow the construction of *immersed walls*. Such a set of immersed walls determines a *wallspace* $(\widetilde{X}, \mathcal{W})$ which yields a dual cube complex $C(\widetilde{X}, \mathcal{W})$ which G acts freely on. Wallspaces were first introduced by Haglund and Paulin [6], and the dual cube complex construction

was first developed by Sageev [8]. In [13] the author defined a criterion, called *dilation*, that determines if an immersed wall produces infinite or finite dimensional cubulations. More precisely, if the immersed walls are *non-dilated*, then $C(\widetilde{X}, \mathcal{W})$ is finite dimensional. We recall the relevant definitions and background in Section 2.

Section 3 establishes a technical result using techniques from [13]. It is shown that that immersed walls can be replaced with *primitive* immersed walls without losing the finite dimensionality or local finiteness of the associated dual cube complex. The reader is encouraged to either read this section alongside [13], or skip it on a first reading.

In Section 4 we analyse $C(\widetilde{X}, \mathcal{W})$ in the finite dimensional case to establish a set of conditions that imply that $G\backslash C(\widetilde{X}, \mathcal{W})$ is virtually special. We decompose $C(\widetilde{X}, \mathcal{W})$ as a tree of spaces, with the same underlying tree as $\widetilde{X}$ and then, under the assumption that the walls are *primitive*, we show that $C(\widetilde{X}, \mathcal{W})$ maps G-equivariantly into $\mathbf{R}^d \times \widetilde{\Gamma}$, where $\mathbf{R}$ is the standard cubulation of $\mathbb{R}$, and $\widetilde{\Gamma}$ is the underlying graph of $\widetilde{X}$. A further criterion, the notion of a *fortified* immersed wall, determines when $C(\widetilde{X}, \mathcal{W})$ is locally finite. Combining these results allow us to give criterion for $G\backslash C(\widetilde{X}, \mathcal{W})$ to be virtually special.

In Section 5 we consider a tubular group acting freely on a CAT(0) cube complex $\widetilde{Y}$. We show that we can obtain from such an action immersed walls that preserve the important properties of $\widetilde{Y}$. More precisely, we prove the following:

Proposition 5.9. *Let G be a tubular group acting freely on a CAT(0) cube complex $\widetilde{Y}$. Then there is a tubular space X and a finite set of immersed walls in X. Moreover, if $(\widetilde{X}, \mathcal{W})$ is the associated wallspace, then*

1. *G acts freely on $C(\widetilde{X}, \mathcal{W})$.*
2. *$C(\widetilde{X}, \mathcal{W})$ is finite dimensional if $\widetilde{Y}$ is finite dimensional.*
3. *$C(\widetilde{X}, \mathcal{W})$ is finite dimensional and locally finite if $\widetilde{Y}$ is locally finite.*

This Proposition is sufficient to allow us to prove 5.10.

In Section 6 we further exploit the results obtained in Section 4 to obtain the following, demonstrating that the cubical dimension of tubular groups with finite dimensional cubulations are virtually within 1 of their cohomological dimension.

Theorem 6.2. *A tubular group acting freely on a finite dimensional CAT(0) cube complex has a finite index subgroup that acts freely on a 3-dimensional CAT(0) cube complex.*

Acknowledgements: I would like to thank Dani Wise and Mark Hagen.

2. Background Tubular Groups and their Cubulations

Let G be a tubular group with associated tubular space X and underlying graph Γ. Given an edge e in a graph we will let $-e$ and $+e$ respectively denote the initial and terminal vertices of e. Let X_v and X_e denote vertex and edge spaces in this graph of spaces. Let X_{-e} and X_{+e} be the boundary circles of X_e, and denote the attaching maps by $\varphi_e^- : X_e^- \to X_{-e}$, and $\varphi_e^+ : X_e^+ \to X_{+e}$. Note that φ_e^- and φ_e^+ respectively represent generators of G_e in G_{e^-} and G_{e^+}. We will let $\widetilde{X}$ denote the universal cover of X. Let $\widetilde{X}_{\widetilde{v}}$ and $\widetilde{X}_{\widetilde{e}}$ denote vertex and edge spaces in the universal cover $\widetilde{X}$, and let $\widetilde{\Gamma}$ denote the Bass-Serre tree. We will assume that each vertex space has the structure of a nonpositively curved geodesic metric space and that attaching maps φ_e^- and φ_e^+ define locally geodesic curves in X_{-e} and X_{+e}.

2.1. **Equitable Sets and Intersection Numbers.** Given a pair of closed curves in a torus $\alpha, \beta : S^1 \to T$, the *intersection points* are the elements $(p, q) \in S^1 \times S^1$ such that $\alpha(p) = \beta(q)$. For a pair of homotopy classes $[\alpha], [\beta]$ of closed curves in a torus T, their *geometric intersection number* $\#\big[[\alpha], [\beta]\big]$ is the minimal number of intersection points realised by a pair of representatives from the respective classes. This number is realised by any pair of geodesic representatives of the classes. If $B = \{[\beta_i]\}$ is a finite set of homotopy classes of curves in T, then $\#[\alpha, B] := \sum_i \#[\alpha, \beta_i]$. Viewing $[\alpha]$ and $[\beta]$ as elements of $\pi_1 T = \mathbb{Z}^2$, we can compute that $\#\big[[\alpha], [\beta]\big] = \det\big[[\alpha], [\beta]\big]$. Given an identification of $\mathbb{Z}^2$ with $\pi_1 T$, the elements of $\mathbb{Z}^2$ are identified with homotopy classes of curves in T, so it makes sense to consider their geometric intersection number. An *equitable set* for a tubular group G is a collection of sets $\{S_v\}_{v \in \Gamma}$, where S_v is a finite set of distinct geodesic curves in X_v disjoint from the attaching maps of adjacent edge spaces, such that S_v generate a finite index subgroup of $\pi_1 X_v = G_v$, and $\#\big[\varphi_e^-, S_{e^-}\big] = \#\big[\varphi_e^+, S_{e^+}\big]$. Note that equitable sets can also be given with S_v a finite subset of G_v that generates a finite index subgroup of G_v and satisfies the corresponding equality for intersection numbers. This is how Wise formulates equitable sets, and its equivalence follows from exchanging elements of $G_v = \pi_1 X_v$ with geodesic closed curves in X_v that represent the corresponding elements. An equitable set is *fortified* if for each edge e in Γ, there exists $\alpha_e^+ \in S_{+e}$ and $\alpha_e^- \in S_{-e}$ such that $\#[\alpha_e^+, \varphi_e^+] = \#[\alpha_e^-, \varphi_e^-] = 0$. An equitable set is *primitive* if every element $\alpha \in S_v$ represents a primitive element in G_v.

2.2. **Immersed Walls From Equitable Sets.** Immersed walls are constructed from *circles* and *arcs*. For each $\alpha \in S_v$, let S^1_α be the domain of α. The disjoint union $\bigsqcup S^1_\alpha$ over all $\alpha \in S_v$ and $v \in V\Gamma$ are the *circles*. Since $\#[\varphi_e^-, S_{e^-}] = \#[\varphi_e^+, S_{e^+}]$,

there exists a bijection from the intersection points between curves in S_{-e} and φ_e^-, and the intersection points between curves in S_{+e} and φ_e^+. Let $(p^-, q^-) \in S^1_{\alpha^-} \times X_e^-$ and $(p^+, q^+) \in S^1_{\alpha^+}$ be corresponding intersection points between $\alpha^\pm \in S_{\pm e}$ and $\varphi_e^\pm$. Then an *arc* $a \cong [0,1]$ has its endpoints attached to p^- and p^+. The endpoints of a are mapped into $X_{-e} \cap X_e$ and $X_{+e} \cap X_e$, so the interior of a can be embedded in X_e. After attaching an arc for each pair of corresponding intersection points, we obtain a set $\{\Lambda_1, \dots, \Lambda_n\}$ of connected graphs that map into X, called *immersed walls*. Each graph Λ_i has its own graph of groups structure with infinite cyclic vertex groups and trivial edge groups. **As in** [13]**, all "immersed walls" in this paper are immersed walls constructed from equitable sets as above.** This means that we are free use the results obtained in [13].

A lift of $\widetilde{\Lambda}_i \to \widetilde{X}$ is a two sided embedding in $\widetilde{X}$, separating $\widetilde{X}$ into two *halfspaces*. The images of the lifts of $\widetilde{\Lambda}_i$ to $\widetilde{X}$ are *horizontal walls* $\mathcal{W}_\mathsf{h}$. The *vertical* walls $\mathcal{W}_\mathsf{v}$ are obtained from the lifts of curves $\alpha_e : S^1 \to X_e$ given by the inclusion $S^1 \times \{0\} \hookrightarrow S^1 \times [-1,1]$. The set $\mathcal{W} = \mathcal{W}_\mathsf{h} \sqcup \mathcal{W}_\mathsf{v}$ of all horizonal and vertical walls gives a *wallspace* $(\widetilde{X}, \mathcal{W})$ where the G-action on $\widetilde{X}$ also gives an action on $\mathcal{W}$. The main theorem of [11] is that a tubular group acts freely on a CAT(0) cube complex if and only if there exists an equitable set.

A set of immersed walls is *fortified* if they are obtained from a fortified equitable set. A set of immersed walls is *primitive* if they are obtained from a primitive equitable set.

Let $\widetilde{\Lambda}$ and $\widetilde{\Lambda}'$ be horizontal walls in $\widetilde{X}$. An point $x \in \widetilde{\Lambda} \cap \widetilde{\Lambda}$ is a *regular intersection point* if it lies in a vertex space $\widetilde{X}_{\widetilde{v}}$, and the lines $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{v}}$ and $\widetilde{\Lambda}' \cap \widetilde{X}_{\widetilde{v}}$ are non-parallel. Otherwise, a point $x \in \widetilde{\Lambda} \cap \widetilde{\Lambda}'$ is a *non-regular intersection point*, and either $x \in \widetilde{X}_{\widetilde{v}}$ where $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{v}} = \widetilde{\Lambda}' \cap \widetilde{X}_{\widetilde{v}}$, or $x \in \widetilde{X}_{\widetilde{e}}$.

An *infinite cube* in a CAT(0) cube complex $\widetilde{Y}$, is an sequence of $c_0, c_1, \dots, c_n, \dots$ such that c_n is an n-cube in $\widetilde{Y}$, and c_n is a face of c_{n+1}. In [13] a *dilation function* is constructed for each immersed wall $R : \pi_1\Lambda \to \mathbb{Q}^*$, and an immersed wall is said to be *dilated* if R has infinite image. The following is Thm 1.2 from that paper:

Theorem 2.1. *Let X be tubular space, and $(\widetilde{X}, \mathcal{W})$ the wallspace obtained from a finite set of immersed walls in X. The following are equivalent:*

1. *The dual cube complex $C(\widetilde{X}, \mathcal{W})$ is infinite dimensional.*
2. *The dual cube complex $C(\widetilde{X}, \mathcal{W})$ contains an infinite cube.*
3. *One of the immersed walls is dilated.*

The following result is also obtained from [13] by combining Thm 1.2, Prop 4.6, and Prop 3.4. The last part follows from the last paragraph of the proof of Prop 3.4.

Proposition 2.2. *Let X be tubular space, and $(\widetilde{X}, \mathcal{W})$ the wallspace obtained from a finite set of immersed walls in X. If $C(\widetilde{X}, \mathcal{W})$ is infinite dimensional, then $\mathcal{W}$ contains an set of pairwise regularly intersecting walls of infinite cardinality, that correspond to the hyperplanes in an infinite cube in $C(\widetilde{X}, \mathcal{W})$. Moreover, the infinite cube contains a canonical 0-cube.*

3. Primitive Immersed Walls

The following result uses the techniques in Section 5. of [13] to compute the dilation function. Let Λ be an immersed wall in X, and let $R : \pi_1(\Lambda) \to \mathbb{Q}^*$ be its dilation function. If R has finite image then Λ is non-dilated. Let $q : \Lambda \to \Omega$ be the quotient map obtained by crushing each circle to a vertex. Note that the arcs in Ω correspond to the arcs in Λ. The dilation function R factors through $q_* : \pi_1\Lambda \to \pi_1\Omega$, so there exists a function $\hat{R} : \pi_1\Omega \to \mathbb{Q}^*$ such that $R = \hat{R} \circ q_*$. We can therefore determine if Λ is dilated by computing the function $\hat{R}$.

We orient each arc in Λ so that all arcs embedded in the same edge space are oriented in the same direction. We orient the arcs in Ω accordingly. We define a weighting $\omega : E(\Omega) \to \mathbb{Q}^*$. Let X_e be an edge space in X, and let σ be an arc mapped into X_e connecting the circles C^- and C^+. Let $\alpha^- : C^- \to X_{-e}$ and $\alpha^+ : C^+ \to X_{+e}$ be the corresponding elements in the equitable set. Then

$$\omega(\sigma) = \frac{\#[\varphi_e^-, \alpha^-]}{\#[\varphi_e^+, \alpha^+]}.$$

If $\gamma = \sigma_1^{\epsilon_1} \cdots \sigma_n^{\epsilon_n}$ is an edge path in Ω where σ_i is an oriented arc in Ω, and $\epsilon \in \pm 1$, then $\hat{R}(\gamma) = \omega(\sigma_1)^{\epsilon_1} \cdots \omega(\sigma_n)^{\epsilon_n}$.

Lemma 3.1. *Let X be a tubular space, and let $G = \pi_1 X$. Let $\Lambda_1, \dots \Lambda_k$ be a set of immersed walls in X obtained from an equitable set $\{S_v\}_{v \in \Gamma}$. Then there exists a set of primitive immersed walls $\Lambda'_1, \dots, \Lambda'_\ell$ in X obtained from an equitable set $\{S'_v\}_{v \in \Gamma}$. Moreover:*

1. *If $\Lambda_1, \dots, \Lambda_k$ are non-dilated, then so are $\Lambda'_1, \dots \Lambda'_\ell$.*
2. *If $\Lambda_1, \dots \Lambda_k$, are fortified, then so are $\Lambda'_1, \dots \Lambda'_\ell$.*

Proof. Each Λ_i decomposes as the union of disjoint circles, which are the domain of locally geodesic closed paths in the equitable set and arcs. Suppose that $\alpha^n \in S_v$, where $[\alpha] \in \pi_1 X_v = G_v$ is primitive. Let Λ_i be the immersed wall containing the circle $S^1_{\alpha^n}$ corresponding to α^n.

A new equitable set is obtained by replacing α^n in S_v with n locally geodesic curves $\{\alpha_i : S^1_i \to X_v\}_{i=1}^n$ with disjoint images in X_v that are isotopic to α in $\widetilde{X}_v$. This remains an equitable set since $\#[\alpha^n, \gamma] = n\#[\alpha, \gamma] = \sum_{i=1}^n \#[\alpha_i, \gamma]$ for

any locally geodesic curve γ in X_v. New immersed walls are obtained from Λ_i by replacing $S^1_{g^n}$ with $S^1_1, \dots S^1_n$ and reattaching the arcs that were attached to the intersection points in $S^1_{\alpha^n}$ to the corresponding intersection points on $S^1_1, \dots S^1_n$. Let $\Lambda_{i1}, \dots, \Lambda_{i\ell}$ be the new set of immersed walls obtained in this way. Note that each arc in $\Lambda_{i1}, \dots, \Lambda_{i\ell}$, corresponds to a unique arc in Λ_i.

Assume that Λ_i is non-dilated. We claim that the new immersed walls $\Lambda_{i1}, \dots, \Lambda_{i\ell}$ are also non-dilated. Let $q_i : \Lambda_i \to \Omega_i$ and $q_{ij} : \Lambda_{ij} \to \Omega_{ij}$ be the quotient maps obtained by crushing the circles to vertices. Let u be the vertex in Ω_i corresponding to $S^1_{\alpha^n}$. Let $R_i : \pi_1\Lambda_i \to \mathbb{Q}^*$ and $R_{ij} : \pi_1\Lambda_{ij} \to \mathbb{Q}^*$ be the dilation functions. Let $\hat{R}_i : \pi_1\Omega_i \to \mathbb{Q}^*$ and $\hat{R}_{ij} : \pi_1\Omega_{ij} \to \mathbb{Q}^*$ be the unique maps such that $R_i = \hat{R}_i \circ q_i$ and $R_{ij} = \hat{R}_{ij} \circ q_{ij}$. Let ω_i and ω_{ij} be the respective weightings of the arcs in Ω_i and Ω_{ij}. By assumption, R_i and $\hat{R}_i$ have finite image. As the arcs in Λ_{ij} correspond to arcs in Λ_i, there is a map $\rho_i : \Omega_{ij} \to \Omega_i$. We show Λ_{ij} is non-dilated by showing that $\hat{R}_{ij}(\gamma) = \hat{R}_i(\rho_i \circ \gamma)$.

Let σ be an oriented arc in Ω_{ij}. The edge $q_{ij}^{-1}(\sigma)$ embeds in an edge space X_e. If the vertices of σ^ϵ are disjoint from $\rho^{-1}(u)$, then $\omega_{ij}(\sigma^\epsilon) = \omega_i(\rho_{ij} \circ \sigma^\epsilon)$. If the endpoints of σ^ϵ are contained in $\rho_{ij}^{-1}(v)$, and correspond to the circles S^1_ι and S^1_τ then

$$\omega_{ij}(\sigma^{\pm 1}) = \frac{\#[\varphi_e^\mp, \alpha_\iota]}{\#[\varphi_e^\pm, \alpha_\tau]} = \frac{n\#[\varphi_e^\mp, \alpha_\iota]}{n\#[\varphi_e^\pm, \alpha_\tau]} = \frac{\#[\varphi_e^\mp, \alpha^n]}{\#[\varphi_e^\pm, \alpha^n]} = \omega(\rho_{ij} \circ \sigma^{\pm 1}).$$

Suppose that exactly one endpoint of σ^ϵ is contained in $\rho_{ij}^{-1}(u)$. If σ^ϵ terminates a vertex in $\rho_{ij}^{-1}(u)$ corresponding to S^1_τ, and the initial vertex corresponds to a circle that is the domain of a locally geodesic curve β then

$$\omega_{ij}(\sigma^{\pm 1}) = \frac{\#[\varphi_e^\mp, \beta]}{\#[\varphi_e^\pm, \alpha_\tau]} = \frac{n\#[\varphi_e^\mp, \beta]}{n\#[\varphi_e^\pm, \alpha_\tau]} = \frac{n\#[\varphi_e^\mp, \beta]}{\#[\varphi_e^\pm, \alpha^n]} = n\omega(\rho_{ij} \circ \sigma^{\pm 1}).$$

If σ^ϵ starts at a vertex in $\rho_{ij}^{-1}(u)$ corresponding to S^1_ι, and the terminal vertex correspond to a circle that is the domain of a locally geodesic curve β then

$$\omega_{ij}(\sigma^{\pm 1}) = \frac{\#[\varphi_e^\mp, \alpha_\iota]}{\#[\varphi_e^\pm, \beta]} = \frac{n\#[\varphi_e^\mp, \alpha_\iota]}{n\#[\varphi_e^\pm, \beta]} = \frac{\#[\varphi_e^\mp, \alpha^n]}{n\#[\varphi_e^\pm, \beta]} = \frac{1}{n}\omega(\rho_{ij} \circ \sigma^{\pm 1}).$$

Therefore, given an edge path γ in Ω_{ij}, since the number of edges exiting vertices in $\rho_{ij}^{-1}(v)$ is the same as the number of vertices entering, $\hat{R}_{ij}(\gamma) = \hat{R}_i(\rho_i \circ \gamma)$.

This procedure produces immersed walls with one fewer non-primitive element in the equitable set. Repeating this procedure for each non-primitive element in the equitable set produces a primitive set of non-dilated immersed walls. It is also clear, that if $\Lambda_1, \dots, \Lambda_k$ are fortified, then so are the new immersed walls. $\square$

4. Finite Dimensional Dual Cube Complexes

Let X be a tubular space and let $G = \pi_1 X$. Let $(\widetilde{X}, \mathcal{W})$ be the wallspace obtained from a set of non-dilated immersed walls $\Lambda_1, \dots \Lambda_k$ constructed from an equitable set, and a vertical immersed wall in each edge space. **We emphasize that in this section all immersed walls are assumed to be non-dilated, even when it is not explicitly stated.** Let $\widetilde{Z} = C(\widetilde{X}, \mathcal{W})$ and let $Z = G\backslash\widetilde{Z}$. By Theorem 2.1, the immersed walls being non-dilated is equivalent to $\widetilde{Z}$ being finite dimensional. For each edge $\widetilde{e}$ in $\widetilde{\Gamma}$ let $\widetilde{\Lambda}_{\widetilde{e}}$ denote the vertical wall in $\widetilde{X}_{\widetilde{e}}$.

We refer to [12] for full background on the dual cube complex construction. A 0-cube z in $\widetilde{Z}$ is a choice of halfspace $z[\widetilde{\Lambda}]$ of $\widetilde{\Lambda}$ for each $\widetilde{\Lambda} \in \mathcal{W}$ such that

(1) If $\widetilde{\Lambda}_1, \widetilde{\Lambda}_2 \in \mathcal{W}$, then $z[\widetilde{\Lambda}_1] \cap z[\widetilde{\Lambda}_2] \neq \emptyset$.
(2) If $x \in X$, then there are only finitely many $\widetilde{\Lambda} \in \mathcal{W}$ such that $x \notin z[\widetilde{\Lambda}]$.

Two 0-cubes z_1, z_2 are adjacent if $z_1[\widetilde{\Lambda}] = z_2[\widetilde{\Lambda}]$ for all but precisely one hyperplane $\widetilde{\Lambda}_{12} \in \mathcal{W}$. The 1-cube joining z_1 and z_2 is dual to the hyperplane corresponding to $\widetilde{\Lambda}_{12}$. An n-cube is then present wherever the 1-skeleton of an n-cube appears. We say that two disjoint walls $\widetilde{\Lambda}, \widetilde{\Lambda}'$ *face* each other in z if $z[\widetilde{\Lambda}]$ is not contained in $z[\widetilde{\Lambda}']$ and vice versa.

Proposition 4.1. *There is a G-equivariant map $f : \widetilde{Z} \to \widetilde{\Gamma}$. Therefore $\widetilde{Z}$ decomposes as a tree of spaces with $\widetilde{Z}_{\widetilde{v}} = f^{-1}(\widetilde{v})$, and $\widetilde{Z}_{\widetilde{e}} = f^{-1}(\widetilde{e})$ is the carrier of the hyperplane corresponding to $\widetilde{\Lambda}_{\widetilde{e}} \in \mathcal{W}_{\mathsf{v}}$.*

By $f^{-1}(e)$ we mean the union of all cubes c in $\widetilde{Z}$ such that $f(c) = e$.

Proof. As there is a vertical wall in each edge space, and since the vertical walls are all disjoint we can identify $C(\widetilde{X}, \mathcal{W}_v)$ with the Bass-Serre tree $\widetilde{\Gamma}$ of $\widetilde{X}$. We define a map $f : \widetilde{Z} \to \widetilde{\Gamma}$: let z be a 0-cube in Z, then define $f(z)$ by letting $f(z)[\widetilde{\Lambda}_e] = z[\widetilde{\Lambda}_e]$. If z_1, z_2 are adjacent 0-cubes, then $z_1[\widetilde{\Lambda}] \neq z_2[\widetilde{\Lambda}]$ for precisely one wall $\widetilde{\Lambda} \in \mathcal{W}$. If $\widetilde{\Lambda}$ is a horizontal wall then $f(z_1) = f(z_2)$ and the 1-cube joining them is also mapped to the same vertex. If $\widetilde{\Lambda} = \widetilde{\Lambda}_e \in \mathcal{W}_{\mathsf{v}}$ then $f(z_1)$ and $f(z_2)$ are adjacent in $\widetilde{\Gamma}$, so the 1-cube joining z_1 and z_2 maps to the edge joining $f(z_1)$ and $f(z_2)$. As f is defined on the 1-skeleton, the map extends uniquely to the entire cube complex $\widetilde{Z}$. Then $\widetilde{Z}_{\widetilde{v}} = f^{-1}(\widetilde{v})$ and $\widetilde{Z}_{\widetilde{e}} = f^{-1}(\widetilde{e})$ is the carrier of the hyperplane corresponding to $\widetilde{\Lambda}_e$. □

Proposition 4.1 implies that Z decomposes as a graph of spaces with vertex spaces Z_v, edge spaces Z_e, and underlying graph $\Gamma = G\backslash\widetilde{\Gamma}$.

The following proposition, which collects the principal consequences of finite dimensionality, is Prop 4.12 in [13].

Proposition 4.2. *Let X be tubular space with geodesic attaching maps, and let $(\widetilde{X}, \mathcal{W})$ be the wallspace obtained from a finite set of immersed walls in X. If the dual cube complex $C(\widetilde{X}, \mathcal{W})$ is finite dimensional, then the horizontal walls in $\mathcal{W}$ can be partitioned into a collection $\mathcal{P}$ of subsets such that:*

1. *The partition $\mathcal{P}$ is preserved by G,*
2. *For each $A \in \mathcal{P}$, the walls in A are pairwise non-intersecting.*
3. *Let $\widetilde{\Lambda} \in A \in \mathcal{P}$ be a wall intersecting $\widetilde{X}_{\widetilde{v}}$. There exists $h \in G_{\widetilde{v}}$ stabilizing an axis in $\widetilde{X}_{\widetilde{v}}$ perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{v}}$ such that $A = \{h^r \widetilde{\Lambda}\}_{r \in \mathbb{Z}}$.*

Any partition of the horizontal walls in $\mathcal{W}$ satisfying conditions (1)-(3) in Proposition 4.2 will be called a *stable partition*.

Lemma 4.3. *Let X be a tubular space and $(\widetilde{X}, \mathcal{W})$ be the wallspace obtained from a finite set of immersed walls in X. Let $\mathcal{P}$ be a stable partition of the horizontal walls in $\mathcal{W}$. Then for each $\widetilde{v} \in \widetilde{\Gamma}$, only finitely many $A \in \mathcal{P}$ contain walls intersecting $\widetilde{X}_{\widetilde{v}}$.*

Proof. Suppose that $\widetilde{\Lambda}$ is a wall intersecting $\widetilde{X}_{\widetilde{v}}$, then, by condition (3) of a stable partition, there exists some $h \in G_{\widetilde{v}}$ that is perpendicular to $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{v}}$ such that $\{h^r \widetilde{\Lambda}\}_{r \in \mathbb{Z}} \in \mathcal{P}$. By G-invariance we can deduce that each of the $G_{\widetilde{v}}$-translates of $\{h^r \widetilde{\Lambda}\}_{r \in \mathbb{Z}}$ is also in $\mathcal{P}$. There are only finitely many such translates, therefore each $G_{\widetilde{v}}$-orbit of a wall in $\widetilde{X}_{\widetilde{v}}$ is contained in finitely many elements of $\mathcal{P}$. The claim then follows from the fact that there are only finitely many $G_{\widetilde{v}}$-orbits of walls intersecting $\widetilde{X}_{\widetilde{v}}$. □

The immersed walls $\Lambda_1, \ldots, \Lambda_k$ are non-dilated, and therefore $\widetilde{Z}$ is finite dimensional, so by Proposition 4.2 there exists a stable partition $\mathcal{P}$ of the horizontal walls in $\mathcal{W}$. Let $\mathcal{P}_{\widetilde{v}}$ be the subpartition containing walls intersecting $\widetilde{X}_{\widetilde{v}}$. Let $\mathcal{P}_{\widetilde{e}}$ be the subpartition of walls intersecting $\widetilde{X}_{\widetilde{e}}$. By Lemma 4.3, both $\mathcal{P}_{\widetilde{v}}$ and $\mathcal{P}_{\widetilde{e}}$ are finite subpartitions. If $\widetilde{e}$ is incident to the vertex $\widetilde{v}$, then $\mathcal{P}_{\widetilde{e}} \subseteq \mathcal{P}_{\widetilde{v}}$. Let $\mathcal{P}_{\widetilde{v}} = \{A_1, \ldots, A_{d_{\widetilde{v}}}\}$. By criterion (3) of a stable partition $A_i = \{h_i^r \widetilde{\Lambda}_i\}_{r \in \mathbb{Z}}$ such that $h_i \in G_{\widetilde{v}}$ stablizes an axis in $\widetilde{X}_{\widetilde{v}}$ perpendicular to $\widetilde{\Lambda}_i \cap \widetilde{X}_{\widetilde{v}}$. The action of $G_{\widetilde{v}}$ preserves both the partition $\mathcal{P}_{\widetilde{v}}$ and the ordering of the walls in each A_i.

Let $\mathbf{R}$ denote the cubulation of $\mathbb{R}$ with a vertex for each integer and an edge joining consecutive integers. Therefore, each 0-cube in $\mathbf{R}^d$ is an element of $\mathbb{Z}^d$. We construct a free action of $G_{\widetilde{v}}$ on $\mathbf{R}^{d_{\widetilde{v}}}$. Let $g \in G_{\widetilde{v}}$ and let $(\alpha_1, \ldots, \alpha_{d_{\widetilde{v}}})$ be a 0-cube in $\mathbf{R}^{d_{\widetilde{v}}}$. Define the map $g \cdot (\alpha_1, \ldots, \alpha_{d_{\widetilde{v}}}) = (\beta_1, \ldots, \beta_{d_{\widetilde{v}}})$ such that $g \cdot h_i^{\alpha_i} \widetilde{\Lambda}_i = h_j^{\beta_j} \widetilde{\Lambda}_j$. As g permutes the walls in $\mathcal{P}_{\widetilde{v}}$, the map g is a bijection on the 0-cubes in $\mathbf{R}^{d_{\widetilde{v}}}$. If $g \cdot h_i^{\alpha_i} \widetilde{\Lambda}_i = h_j^{\beta_j} \widetilde{\Lambda}_j$, then necessarily $g \cdot h_i^{\alpha_i + 1} \widetilde{\Lambda}_i = h_j^{\beta_j \pm 1} \widetilde{\Lambda}_j$, so adjacent 0-cubes are mapped to adjacent 0-cubes and the map extends to an isomorphism of $\mathbf{R}^{d_{\widetilde{v}}}$. If

$g \cdot (\alpha_1, \dots, \alpha_{d_{\widetilde{v}}}) = (\alpha_1, \dots, \alpha_{d_{\widetilde{v}}})$ then g would stabilize all the walls in $\mathcal{P}_{\widetilde{v}}$, which would imply that it fixed every 0-cube in $\widetilde{Z}_{\widetilde{v}}$. Since $G_{\widetilde{v}}$ acts freely on $\widetilde{Z}_{\widetilde{v}}$ this would imply that $g = 1_G$, and hence $G_{\widetilde{v}}$ acts freely on $\mathbf{R}^{d_{\widetilde{v}}}$.

We also define an embedding $\phi_{\widetilde{v}} : \widetilde{Z}_{\widetilde{v}} \to \mathbf{R}^{d_{\widetilde{v}}}$. If z is a 0-cube in $\widetilde{Z}_{\widetilde{v}}$ then every wall $\widetilde{\Lambda}$ that is either vertical or not in contained in the subpartition $\mathcal{P}_{\widetilde{v}}$ has $\widetilde{X}_{\widetilde{v}} \subseteq z[\widetilde{\Lambda}]$. Therefore z is entirely determined by $z[\widetilde{\Lambda}]$ for $\widetilde{\Lambda}$ in $\mathcal{P}_{\widetilde{v}}$. For $1 \leq i \leq d\widetilde{v}$ the set $\{h_i^r \widetilde{\Lambda}_i \cap \widetilde{X}_{\widetilde{v}}\}_{r \in \mathbb{Z}}$ is an infinite collection of disjoint parallel lines in $\widetilde{X}_{\widetilde{v}}$. As all the walls in A_i are disjoint in $\widetilde{X}$, for each 0-cube z in $\widetilde{Z}_{\widetilde{v}}$ there exists a unique $\alpha_i \in \mathbb{Z}$ such that $h_i^{\alpha_i}\widetilde{\Lambda}_i$ and $h_i^{\alpha_i+1}\widetilde{\Lambda}_i$ face each other in z. Let $\phi_{\widetilde{v}}(z) = (\alpha_1, \dots, \alpha_{d_{\widetilde{v}}})$. Note that the map is injective and sends adjacent 0-cubes to adjacent 0-cubes, so the map on the 0-cubes extends to an embedding of the entire cube complex.

Lemma 4.4. *The embedding $\phi_{\widetilde{v}} : \widetilde{Z}_{\widetilde{v}} \to \mathbf{R}^{d_{\widetilde{v}}}$ is $G_{\widetilde{v}}$-equivariant.*

Proof. Let $g \in G_{\widetilde{v}}$. If $\phi_{\widetilde{v}}(z) = (\alpha_1, \dots, \alpha_{d_{\widetilde{v}}})$ and $g \cdot \phi_{\widetilde{v}}(z) = (\beta_1, \dots, \beta_{d_{\widetilde{v}}})$, then $(gz)[h_i^{\alpha_i}\widetilde{\Lambda}_i] = z[gh_i^{\alpha_i}\widetilde{\Lambda}_i] = z[h_j^{\beta_j}\widetilde{\Lambda}_j]$, which implies that $\phi_{\widetilde{v}}(gz) = (\beta_1, \dots, \beta_{d_{\widetilde{v}}}) = g\phi_{\widetilde{v}}(z)$. □

Let $\widetilde{e}$ be an edge adjacent to $\widetilde{v}$. Then either $+\widetilde{e} = \widetilde{v}$ or $-\widetilde{e} = \widetilde{v}$. We define a free action of $G_{\widetilde{e}}$ on $\mathbf{R}^{d_{\widetilde{e}}} \times [-1, 1]$. After reindexing, let $\mathcal{P}_{\widetilde{e}} = \{A_1, \dots, A_{d_{\widetilde{e}}}\} \subseteq \mathcal{P}_{\widetilde{v}}$ where $d_{\widetilde{e}} \leq d_{\widetilde{v}}$. Let $(\alpha_1, \dots, \alpha_{d_{\widetilde{e}}}, \pm 1)$ be a 0-cube in $\mathbf{R}^d \times [-1, 1]$ and let $g \in G_{\widetilde{e}}$. Then $g \cdot (\alpha_1, \dots, \alpha_{d_{\widetilde{e}}}, \pm 1) = (\beta_1, \dots, \beta_{d_{\widetilde{e}}}, \pm 1)$ such that $g \cdot h_i^{\alpha_i}\widetilde{\Lambda}_i = h_j^{\beta_j}\widetilde{\Lambda}_j$. As in the case of vertex spaces, this map extends to an isomorphism of $\mathbf{R}^{d_{\widetilde{e}}} \times [-1, 1]$.

As with the vertex spaces, there is a $G_{\widetilde{e}}$-equivariant embedding $\phi_{\widetilde{e}} : \widetilde{Z}_{\widetilde{e}} \to \mathbf{R}^{d_{\widetilde{e}}} \times [-1, 1]$. Let z be a 0-cube in $\widetilde{Z}_{\widetilde{e}}$. Then for each $1 \leq i \leq d_{\widetilde{e}}$ there exists a unique α_i such that $h_i^{\alpha_i}\widetilde{\Lambda}_i$ faces $h_i^{\alpha_i+1}\widetilde{\Lambda}_i$ in z, and $\widetilde{X}_{\pm\widetilde{e}} \subseteq z[\widetilde{\Lambda}_{\widetilde{e}}]$. Define $\phi_{\widetilde{e}}(z) = (\alpha_1, \dots, \alpha_{d_{\widetilde{e}}}, \pm 1)$.

Let $\widetilde{v} = \pm\widetilde{e}$. The free action of $G_{\widetilde{v}}$ on $\mathbf{R}^{d_{\widetilde{v}}}$ restricts to a free action of $G_{\widetilde{e}}$. We claim that we can embed $\mathbf{R}^{d_{\widetilde{e}}} \times \{\pm 1\}$ into $\mathbf{R}^{d_{\widetilde{v}}}$ in a $G_{\widetilde{e}}$-equivariant way. Let $H_{\widetilde{e}} \subseteq \widetilde{Z}$ be the hyperplane corresponding to $\widetilde{\Lambda}_e$. As $\widetilde{Z}_{\widetilde{e}}$ is the carrier of $H_{\widetilde{e}}$, we can identify $\widetilde{Z}_{\widetilde{e}}$ with $H_{\widetilde{e}} \times [-1, 1]$. Note that $H_{\widetilde{e}} \times \{\pm 1\}$ embeds as a subspace in $\widetilde{Z}_{\widetilde{v}}$, and $\phi_{\widetilde{e}}$ restricts to an embedding $\phi_{\widetilde{e}}^{\pm} : H_{\widetilde{e}} \times \{\pm 1\} \to \mathbf{R}^{d_{\widetilde{e}}}$, where $\widetilde{v} = \pm\widetilde{e}$.

We construct an embedding $\Psi_{\widetilde{e}}^{\pm} : \mathbf{R}^{d_{\widetilde{e}}} \to \mathbf{R}^{d_{\widetilde{v}}}$. Recall that $\mathcal{P}_{\widetilde{e}} = \{A_1, \dots, A_{d_{\widetilde{e}}}\} \subseteq \mathcal{P}_{\widetilde{v}} = \{A_1, \dots, A_{d_{\widetilde{v}}}\}$. For $d_{\widetilde{e}} < j \leq d_{\widetilde{v}}$ if $h_j^r\widetilde{\Lambda}_j \in A_j$, then $\widetilde{X}_{\widetilde{e}} \subseteq z[h_j^r\widetilde{\Lambda}_j]$ for all 0-cubes z in $\widetilde{Z}_{\widetilde{e}}$. Therefore, there is a unique $\alpha_j^{\widetilde{e}} \in \mathbb{Z}$ such that $h_j^{\alpha_j^{\widetilde{e}}}\widetilde{\Lambda}_j$ faces $h_j^{\alpha_j^{\widetilde{e}}+1}\widetilde{\Lambda}_j$ for every 0-cube z in $\widetilde{Z}_{\widetilde{e}}$ and $d_{\widetilde{e}} < j \leq d_{\widetilde{v}}$. Thus we define

$$\Psi_{\widetilde{e}}^{\pm}(\alpha_1, \dots, \alpha_{d_{\widetilde{e}}}) = (\alpha_1, \dots, \alpha_{d_{\widetilde{e}}}, \alpha_{d_{\widetilde{e}}+1}^{\widetilde{e}}, \dots, \alpha_{d_{\widetilde{v}}}^{\widetilde{e}}).$$

The $G_{\widetilde{e}}$-equivariance of $\Psi_{\widetilde{e}}^{\pm}$ will require a further assumption:

Lemma 4.5. *The following commutative square is $G_{\widetilde{e}}$-equivariant provided the immersed walls are primitive.*

$$\begin{array}{ccc} H_{\widetilde{e}} \times \{\pm 1\} & \xrightarrow{\phi^{\pm}_{\widetilde{e}}} & \mathbf{R}^{d_{\widetilde{e}}} \\ \downarrow & & \downarrow \Psi^{\pm}_{\widetilde{e}} \\ \widetilde{Z}_{\widetilde{v}} & \xrightarrow{\phi_{\widetilde{v}}} & \mathbf{R}^{d_{\widetilde{v}}} \end{array}$$

Moreover, $\Psi^{\pm}_{\widetilde{e}}$ is a $G_{\widetilde{e}}$-equivariant inclusion that is equivalent to extending the $G_{\widetilde{e}}$-action on $\mathbf{R}^{d_{\widetilde{e}}}$ by a trivial action on $\mathbf{R}^{d_{\widetilde{v}}-d_{\widetilde{e}}}$.

Proof. Let z be a 0-cube in $H_{\widetilde{e}} \times \{\pm 1\}$. Then by construction

$$\Psi^{\pm}_{\widetilde{e}} \circ \phi^{\pm}_{\widetilde{e}}(z) = (\alpha_1, \dots, \alpha_{d_{\widetilde{e}}}, \alpha^{\widetilde{e}}_{d_{\widetilde{e}}+1}, \dots, \alpha^{\widetilde{e}}_{d_{\widetilde{v}}}) = \phi_{\widetilde{v}}(z).$$

To verify that $\Psi_{\widetilde{e}}$ is $G_{\widetilde{e}}$-equivariant, let $g \in G_{\widetilde{e}}$. For $1 \leq i \leq d_{\widetilde{e}}$ there exists $1 \leq j \leq d_{\widetilde{e}}$ and β_j be such that $g \cdot h_i^{\alpha_i} \widetilde{\Lambda}_i = h_j^{\beta_j} \widetilde{\Lambda}_j$. For $d_{\widetilde{e}} < i \leq d_{\widetilde{v}}$ the intersection $\widetilde{\Lambda}_i \cap \widetilde{X}_{\widetilde{v}}$ is a geodesic line parallel to $\widetilde{X}_{\widetilde{e}} \cap \widetilde{X}_{\widetilde{v}}$. Thus, $G_{\widetilde{e}}$ stabilizes $\widetilde{\Lambda}_i \cap \widetilde{X}_{\widetilde{v}}$. As the immersed walls are primitive we can deduce that $G_{\widetilde{e}}$ stabilizes $\widetilde{\Lambda}_i$. For $d_{\widetilde{e}} < i \leq d_{\widetilde{v}}$ we deduce that $g \cdot h_i^{\alpha} \widetilde{\Lambda}_i = h_i^{\alpha} \widetilde{\Lambda}_i$ for $\alpha \in \mathbb{Z}$, and conclude:

$$\begin{aligned} g \cdot \Psi_{\widetilde{e}}(\alpha_1, \dots, \alpha_{\widetilde{e}}) &= g \cdot (\alpha_1, \dots, \alpha_{d_{\widetilde{e}}}, \alpha^{\widetilde{e}}_{d_{\widetilde{e}}+1}, \dots, \alpha^{\widetilde{e}}_{d_{\widetilde{v}}}) \\ &= (\beta_1, \dots, \beta_{d_{\widetilde{e}}}, \alpha^{\widetilde{e}}_{d_{\widetilde{e}}+1}, \dots, \alpha^{\widetilde{e}}_{d_{\widetilde{v}}}) \\ &= \Psi_{\widetilde{e}}(\beta_1, \dots, \beta_{\widetilde{e}}) \\ &= \Psi_{\widetilde{e}}(g \cdot (\alpha_1, \dots, \alpha_{\widetilde{e}})). \end{aligned}$$

Observe that $G_{\widetilde{e}}$ acts trivially on the last $d_{\widetilde{v}} - d_{\widetilde{e}}$ coordinates. □

Let $d = \max\{|\mathcal{P}_{\widetilde{v}}| \mid \widetilde{v} \in V\widetilde{\Gamma}\}$, which is finite, since there are only finitely many vertex orbits.

Proposition 4.6. *If the immersed walls $\Lambda_1, \dots, \Lambda_k$ are primitive, then G acts freely on $\mathbf{R}^d \times \widetilde{\Gamma}$ such that the action on the $\widetilde{\Gamma}$ factor is the action of G on the Bass-Serre tree. Moreover, there is a G-equivariant embedding $\phi : \widetilde{Z} \to \mathbf{R}^d \times \widetilde{\Gamma}$.*

Proof. The $G_{\widetilde{v}}$ and $G_{\widetilde{e}}$-actions on $\mathbf{R}^{d_{\widetilde{v}}}$ and $\mathbf{R}^{d_{\widetilde{e}}}$ can be equivariantly extended to actions on $\mathbf{R}^d$ such that $G_{\widetilde{v}}$ and $G_{\widetilde{e}}$ act trivially on the additional factors. Therefore, the $G_{\widetilde{e}}$-commutative square in Lemma 4.5 can be extended:

$$\begin{array}{ccccc} H_{\widetilde{e}} \times \{\pm 1\} & \xrightarrow{\phi^{\pm}_{\widetilde{e}}} & \mathbf{R}^{d_{\widetilde{e}}} & \longrightarrow & \mathbf{R}^{d} \\ \downarrow & & \downarrow \Psi^{\pm}_{\widetilde{e}} & & \downarrow \\ \widetilde{Z}_{\widetilde{v}} & \xrightarrow{\phi_{\widetilde{v}}} & \mathbf{R}^{d_{\widetilde{v}}} & \longrightarrow & \mathbf{R}^{d} \end{array}$$

Therefore, we obtain a G-action on $\mathbf{R}^d \times \widetilde{\Gamma}$ and a G-equivariant embedding of the tree of spaces $\widetilde{Z} \to \mathbf{R}^d \times \widetilde{\Gamma}$. □

Proposition 4.7. *$\widetilde{Z}$ is locally finite if and only if $\Lambda_1, \dots, \Lambda_k$ are fortified.*

Proof. If $\Lambda_1, \dots, \Lambda_k$ is not fortified, then there exists a vertex space $\widetilde{X}_{\widetilde{v}}$ and an adjacent edge space $\widetilde{X}_{\widetilde{e}}$ such that every horizontal wall $\widetilde{\Lambda}$ in $\mathcal{P}_{\widetilde{v}}$ intersects $\widetilde{X}_{\widetilde{v}}$ as a line $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{v}}$ that intersects $\widetilde{X}_{\widetilde{v}} \cap \widetilde{X}_{\widetilde{e}}$. Therefore, every horizontal wall intersecting $\widetilde{X}_{\widetilde{v}}$ intersects $\widetilde{X}_{\widetilde{e}}$, so $\mathcal{P}_{\widetilde{e}} = \mathcal{P}_{\widetilde{v}}$. Let $\widetilde{e}_1, \dots, \widetilde{e}_i, \dots$ be an enumeration of the $G_{\widetilde{v}}$-orbit of $\widetilde{e}$. Then $\mathcal{P}_{\widetilde{e}_i} = \mathcal{P}_{\widetilde{v}}$ and $\widetilde{\Lambda}_{\widetilde{e}_i}$ intersects all the horizontal walls in $\mathcal{P}_{\widetilde{v}}$.

Let z be a 0-cube in $\widetilde{Z}_{\widetilde{v}}$. There is a 0-cube z_i such that $z_i[\widetilde{\Lambda}] = z[\widetilde{\Lambda}]$ for $\widetilde{\Lambda} \neq \widetilde{\Lambda}_{\widetilde{e}_i}$, and $z_i[\widetilde{\Lambda}_{\widetilde{e}_i}] \neq z[\widetilde{\Lambda}_{\widetilde{e}_i}]$. To verify z_i is a 0-cube note that every wall in $\mathcal{P}_{\widetilde{v}}$ intersects $\widetilde{\Lambda}_{\widetilde{e}}$, and every other wall $\widetilde{\Lambda}$ that is not $\widetilde{\Lambda}_e$ has $\widetilde{X}_{\widetilde{e}} \subseteq z_i[\widetilde{\Lambda}]$. Therefore $z_i[\widetilde{\Lambda}_{\widetilde{e}}] \cap z_i[\widetilde{\Lambda}] \neq \emptyset$ for all $\widetilde{\Lambda} \in \mathcal{W} - \{\widetilde{\Lambda}_{\widetilde{e}}\}$. For any walls $\widetilde{\Lambda}_1, \widetilde{\Lambda}_2 \in \mathcal{W} - \{\widetilde{\Lambda}_{\widetilde{e}}\}$ the intersection $z_i[\widetilde{\Lambda}_1] \cap z_i[\widetilde{\Lambda}_2] = z[\widetilde{\Lambda}_1] \cap z[\widetilde{\Lambda}_2] \neq \emptyset$. Finally, if $x \in \widetilde{X}$, then $x \in z_i[\widetilde{\Lambda}]$ for all but finitely many $\widetilde{\Lambda} \in \mathcal{W}$, because it is true for z, which differs from z_i on precisely one wall. Each z_i is adjacent to z since they differ on precisely one wall, so $z_1, \dots, z_i, \dots$ is an infinite collection of distinct 0-cubes adjacent to z, and $\widetilde{Z}$ is not locally finite.

To show that converse, we first observe that the embedding $\phi_{\widetilde{v}} : \widetilde{Z}_{\widetilde{v}} \to \mathbf{R}^{d_{\widetilde{v}}}$ proves that $\widetilde{Z}_{\widetilde{v}}$ is always locally finite, irrespective of whether the immersed walls are fortified. Let z be a 0-cube in $\widetilde{Z}_{\widetilde{v}}$, and let $\widetilde{u}$ be adjacent to $\widetilde{v}$ in $\widetilde{\Gamma}$ via an edge $\widetilde{e}$. Then z can be adjacent to at most one 0-cube $z_{\widetilde{e}}$ in $\widetilde{Z}_{\widetilde{u}}$ such that $z[\widetilde{\Lambda}] = z_{\widetilde{e}}[\widetilde{\Lambda}]$ for all $\widetilde{\Lambda} \in \mathcal{W}$ except $\widetilde{\Lambda}_e$. This $z_{\widetilde{e}}$ may not always define a 0-cube however. Let $\widetilde{e}$ be an edge adjacent to $\widetilde{v}$. As the immersed walls are fortified there exists $\{h^r\widetilde{\Lambda}\}_{r\in\mathbb{Z}} \in \mathcal{P}_{\widetilde{v}}$ such that $\{h^r\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{v}}\}_{r\in\mathbb{Z}}$ is an infinite set of lines parallel to $\widetilde{X}_{\widetilde{e}} \cap \widetilde{X}_{\widetilde{v}}$. As $\{h^r\widetilde{\Lambda}\}$ is a set of disjoint walls, there exists r such that $h^r\widetilde{\Lambda}$ and $h^{r+1}\widetilde{\Lambda}$ are facing in z. There are only finitely many edges $g_1\widetilde{e}, \dots, g_m\widetilde{e} \in G_{\widetilde{v}}\widetilde{e}$ such that $\widetilde{X}_{g_i\widetilde{e}} \subseteq z[h^r\widetilde{\Lambda}] \cap z[h^{r+1}\widetilde{\Lambda}]$. If $g\widetilde{e}$ is an edge such that $\widetilde{X}_{g_i\widetilde{e}}$ is not contained in $z[h^r\widetilde{\Lambda}] \cap z[h^{r+1}\widetilde{\Lambda}]$, then either $z_{g\widetilde{e}}[h^r\widetilde{\Lambda}] \cap z_{g\widetilde{e}}[\widetilde{\Lambda}_e] = \emptyset$ or $z_{g\widetilde{e}}[h^{r+1}\widetilde{\Lambda}] \cap z_{g\widetilde{e}}[\widetilde{\Lambda}_e] = \emptyset$ so $z_{g\widetilde{e}}$ is not a 0-cube. As there are only finitely many $G_{\widetilde{v}}$-orbits of edges incident to $\widetilde{v}$ we conclude that $\widetilde{z}_{\widetilde{e}}$ is a 0-cube for finitely many edges $\widetilde{e}$ incident to $\widetilde{v}$. □

Proposition 4.8. *If $\Lambda_1, \dots, \Lambda_k$ are primitive, fortified, non-dilated immersed walls, then G is virtually special.*

Proof. By Proposition 4.6, There is a free action of G on $\mathbf{R}^d \times \mathrm{Aut}([-1,1]^d)$, so G is a subgroup of $\mathrm{Isom}(\mathbf{R}^d \times \widetilde{\Gamma}) \cong \left(\mathbb{Z}^d \rtimes \mathrm{Aut}([-1,1])^d\right) \times \mathrm{Aut}(\widetilde{\Gamma})$. Therefore, there is a projection $\rho : G \to \mathbb{Z}^d \rtimes \mathrm{Aut}([-1,1]^d)$. Each vertex group $G_{\widetilde{v}}$ embeds in $\mathbb{Z}^d$, and the mapping is invariant under conjugation. As there are only finitely orbits of vertices in $\widetilde{\Gamma}$, there exists a finite index subgroup $(D\mathbb{Z})^d \leqslant \mathbb{Z}^d$ such that if $\widetilde{e}$ is incident to $\widetilde{v}$ then

$G_{\widetilde{e}} \cap (D\mathbb{Z})^d$ is generated by a primitive element in $G_{\widetilde{v}} \cap (D\mathbb{Z})^d$. Let $G' = \rho^{-1}\big((D\mathbb{Z})^d\big)$. Then $G' \leqslant G$ is a finite index subgroup that embeds in $(D\mathbb{Z})^d \times \mathrm{Aut}(\widetilde{\Gamma})$ such that each edge group is generated by an element that is primitive in the adjacent vertex groups. By Proposition 4.6 there is a G-equivariant embedding $\phi : \widetilde{Z} \to \mathbf{R}^d \times \widetilde{\Gamma}$. As G' does not permute the factors of $\mathbf{R}^d$ we can deduce that the hyperplanes in $\big(\mathbf{R}^d \times \widetilde{\Gamma}\big)/G$ do not self-intersect, so neither do the hyperplanes in $\widetilde{Z}/G'$. Indeed, they are also 2-sided and cannot inter-osculate.

Let G'' be a finite index subgroup such that the underlying graph Γ'' has girth at least 2. Let $\widetilde{e}$ be an edge such that $+\widetilde{e} = \widetilde{v}$. As $\Lambda_1, \dots, \Lambda_k$ are fortified, we conclude that $d_{\widetilde{v}} > d_{\widetilde{e}}$ and $\widetilde{Z}^+_{\widetilde{e}}$ is a proper subcomplex of $\widetilde{Z}_{\widetilde{v}}$. As $G''_{\widetilde{e}}$ is primitive, if $g \in G''_{\widetilde{v}} - G''_{\widetilde{e}}$ then $g h^r_{d_{\widetilde{v}}} \widetilde{\Lambda}_{d_{\widetilde{v}}} \neq h^r_{d_{\widetilde{v}}} \widetilde{\Lambda}_{d_{\widetilde{v}}}$ as g acts by translation on $\widetilde{X}_{\widetilde{v}}$ in a direction non-parallel to $\widetilde{\Lambda}_{d_{\widetilde{v}}} \cap \widetilde{X}_{\widetilde{v}}$. Thus, $\widetilde{Z}^{\pm}_{\widetilde{e}}$ is not stabilizes by g, so we can deduce that $\mathrm{Stab}_{G''_{\widetilde{v}}}(\widetilde{Z}^+_{\widetilde{e}}) = G''_{\widetilde{e}}$. Therefore, $G''_{\widetilde{e}} \backslash \widetilde{Z}^+_{\widetilde{e}}$ embeds in $G''_{\widetilde{v}} \backslash \widetilde{Z}_{\widetilde{v}}$.

Let $Z'' = G'' \backslash \widetilde{Z}$. Let z be a 0-cube in Z''_v. Let H_e be the vertical hyperplane contained in Z''_e, and dual to an edge incident to v. As the attaching maps of Z''_e are embeddings, and Γ'' has girth at least s, we deduce that z can only be incident to one end of a single 1-cube intersected by H_e. Therefore H_e does not self-osculate.

Let σ be a 1-cube in $\mathbf{R}^d \times \widetilde{\Gamma}$ that projects to a 1-cube in $\mathbf{R}^d$. The G''-orbit of σ is a set of 1-cubes, that all project to the same factor of $\mathbf{R}^d$, since G'' does not permute the factors of $\mathbf{R}^d$. As G'' does not invert hyperplanes, after subdividing $\mathbf{R}^d$ we can assume that the G''-orbit of σ is a disjoint set of 1-cubes. Therefore, after the corresponding subdivision, we conclude that the horizontal hyperplanes in Z'' don't self-osculate. □

We note that the requirement in Proposition 4.8 that the immersed walls are fortified is necessary, as the following example demonstrates.

Example 4.9. Let $G = \langle a, b, t \mid [a, b] = 1, tat^{-1} = a \rangle$. We can decompose G as the cyclic HNN extension of the vertex group $G_v = \langle a, b \rangle$ with stable letter t. Thus, G is a tubular group. Let X be the corresponding tubular space with a single vertex space X_v and edge space X_e. There is an equitable set $\{\alpha_1, \alpha_2\}$ where α_1 is a geodesic curve in X_v representing $ab \in G_v$, and α_2 is a geodesic curve in X_v representing $ab^{-1} \in G_v$. Note that each attaching map φ^+_e and φ^-_e intersects each curve in the equitable set precisely once. Therefore, we obtain a pair of embedded immersed horizontal walls Λ_1 and Λ_2, by connecting respective intersection points with φ^+_e and φ^-_e by an arc. A vertical wall Λ_e is also embedded in X_e.

In the wallspace $(\widetilde{X}, \mathcal{W})$ we can decompose $\mathcal{W}$ into three sets of disjoint walls: the walls $\mathcal{W}_1$ that cover Λ_1, the walls $\mathcal{W}_2$ that cover Λ_2, and the walls $\mathcal{W}_e$ that cover Λ_e. These walls are disjoint since the immersed walls are embedded. Furthermore, the

walls in different sets pairwise intersect. Therefore we can conclude that $C(\widetilde{X}, \mathcal{W}) = \mathbf{R}^2 \times \widetilde{\Gamma}$. As this is not locally finite, $G\backslash C(\widetilde{X}, \mathcal{W})$ cannot be virtually special.

5. Revisiting Equitable Sets

Although Wise proved in [11] that acting freely on a CAT(0) cube complex $\widetilde{Y}$ implied the existence of an equitable set, and thus a system of immersed walls as in Section 2, no relationship was established between $\widetilde{Y}$ and the resulting dual $C(\widetilde{X}, \mathcal{W})$. Proposition 5.9 gives the relationship required to reduce Theorem 5.10 to considering cubulations obtained from equitable sets.

This section will apply the following theorem from [14]. A *cubical quasiline* is a CAT(0) cube complex quasi-isomorphic to $\mathbb{R}$.

Theorem 5.1. *Let G be virtually $\mathbb{Z}^n$. Suppose G acts properly and without inversions on a CAT(0) cube complex $\widetilde{Y}$. Then G stabilizes a finite dimensional subcomplex $\widetilde{Z} \subseteq \widetilde{Y}$ that is isometrically embedded in the combinatorial metric, and $\widetilde{Z} \cong \prod_{i=1}^{m} C_i$, where each C_i is a cubical quasiline and $m \geq n$. Moreover, $\mathrm{Stab}_G(\Lambda)$ is a codimension-1 subgroup for each hyperplane Λ in $\widetilde{Z}$.*

Theorem 5.1 allows us to prove the following:

Lemma 5.2. *Let G be a tubular group acting freely on a CAT(0) cube complex $\widetilde{Y}$. Let G_v be a vertex group in G, then there exists a G_v-equivariant subspace $\widetilde{X}_v \subseteq \widetilde{Y}$ homeomorphic to $\mathbb{R}^2$. Moreover, $\widetilde{X}_{\widetilde{v}}$ has a metric such that the intersection of a hyperplane in $\widetilde{Y}$ with $\widetilde{X}_v$ is either empty, or a geodesic line.*

Proof. By Theorem 5.1, there exists a G_v-equivariant subcomplex $\widetilde{Y}_v \subseteq \widetilde{Y}$ that isometrically embeds in the combinatorial metric, and such that $\widetilde{Y}_v \cong \prod_{i=1}^{m} C_i$ where each C_i is a cubical quasiline. By the flat torus theorem [2], G_v stabilizes a flat $\widetilde{X}_v \subseteq \widetilde{Y}_v$, that is a convex subset in the CAT(0) metric of $\widetilde{Y}_v$. As the stabilizers of hyperplanes in $\widetilde{Y}_v$ are codimension-1 subgroups of G_v, the intersection of a hyperplane in $\widetilde{Y}$ with $\widetilde{X}_v$ is either empty or a geodesic line in the CAT(0) metric inherited from $\widetilde{Y}_v$. □

If S is a subset of a CAT(0) cube complex $\widetilde{Y}$, then let $\mathsf{hull}(S)$ denote the *combinatorial convex hull* of S. The combinatorial convex hull of S is the minimal convex subcomplex containing S. Equivalently, $\mathsf{hull}(S)$ is the intersection of all closed halfspaces containing S.

Definition 5.3. Let X be a tubular space and let Y be a nonpositively curved cube complex. A map $f : X \to Y$ is an *amicable immersion* if:

(1) $f_* : \pi_1 X \to \pi_1 Y$ is an isomorphism.

(2) The G-equivariant map $\widetilde{f} : \widetilde{X} \to \widetilde{Y}$ embeds each vertex space $\widetilde{X}_{\widetilde{v}}$ in $\widetilde{Y}$.
(3) Each $\widetilde{X}_{\widetilde{v}}$ has a Euclidean metric such that if $H \subseteq \widetilde{Y}$ is a hyperplane, then the intersection $H \cap \widetilde{X}_{\widetilde{v}}$ is either the empty set, or a single geodesic line in $\widetilde{X}_{\widetilde{v}}$.
(4) Each edge space $\widetilde{X}_{\widetilde{e}}$ is emdedded transverse to the hyperplanes.
(5) Each $\widetilde{X}_{\widetilde{e}}$ is contained in $\mathsf{hull}\big(\bigcup_{\widetilde{v}\in V\widetilde{\Gamma}} \widetilde{X}_{\widetilde{v}}\big)$.

Note that the Euclidean metric on each $\widetilde{X}_{\widetilde{v}}$ is not the subspace metric induced from $\widetilde{Y}$.

Lemma 5.4. *Let X be a tubular space and let Y be a nonpositively curved cube complex. Let $F : \pi_1 X \to \pi_1 Y$ be an isomorphism. Then there is an amicable immersion $f : X \to Y$ such that $f_* = F$.*

Proof. Use F to identify $G = \pi_1 X$ with $\pi_1 Y$. The claim is proven by constructing a G-equivariant map between the tree of spaces $\widetilde{X} \to \widetilde{Y}$. By Lemma 5.2 for each $\widetilde{v} \in V\widetilde{\Gamma}$, we can $G_{\widetilde{v}}$-equivariantly embed a Euclidean flat $\widetilde{X}_{\widetilde{v}}$ in $\widetilde{Y}$ such that if $H \subseteq \widetilde{Y}$ is a hyperplane, then the intersection $H \cap \widetilde{X}_{\widetilde{v}}$ is either the empty set, or a single geodesic line in $\widetilde{X}_{\widetilde{v}}$. Moreover, we can ensure that $\bigcup_{\widetilde{v}\in V\widetilde{\Gamma}} \widetilde{X}_{\widetilde{v}}$ is G-equivariant. The edges spaces $\widetilde{X}_{\widetilde{e}}$ can then be inserted transverse to the hyperplanes in $\widetilde{Y}$ so that the intersections $\widetilde{X}_{\widetilde{e}} \cap \widetilde{X}_{\widetilde{v}}$ with adjacent vertex spaces is not contained in a hyperplane in $\widetilde{Y}$, and $\widetilde{X}_{\widetilde{e}}$ is contained inside $\mathsf{hull}\big(\bigcup_{\widetilde{v}\in V\widetilde{\Gamma}} \widetilde{X}_{\widetilde{v}}\big)$. □

Lemma 5.5. *Let $X \to Y$ be an amicable immersion, where Y is finite dimensional. If $\widetilde{v}$ is a vertex in $\widetilde{\Gamma}$, then $\mathsf{hull}(\widetilde{X}_{\widetilde{v}})$ embeds as a subcomplex of $\mathbf{R}^d$ for some d.*

Proof. Let $G = \pi_1 X$. Let y be a 0-cube in $\mathsf{hull}(\widetilde{X}_{\widetilde{v}})$. If H is a hyperplane in $\widetilde{Y}$, let $y[H]$ denote the halfspace of H containing y. Each 0-cube is determined by the halfspace containing it for each hyperplane. If H is a hyperplane that doesn't intersect $\widetilde{X}_{\widetilde{v}}$, then $y[H]$ is the halfspace containing $\widetilde{X}_{\widetilde{v}}$, and therefore is $y[H]$ fixed for all 0-cubes y in $\mathsf{hull}(\widetilde{X}_{\widetilde{v}})$.

Let $\mathcal{H}_{\widetilde{v}}$ denote the hyperplanes intersecting $\widetilde{X}_{\widetilde{v}}$. Let $H \in \mathcal{H}_{\widetilde{v}}$. The intersection $\widetilde{X}_{\widetilde{v}} \cap H$ is a geodesic line in $\widetilde{X}_{\widetilde{v}}$. Let $g \in G_{\widetilde{v}}$ be an isometry that stabilizes an axis in $\widetilde{X}_{\widetilde{v}}$ that is not parallel to $\widetilde{X}_{\widetilde{v}} \cap H$. Then $\{g^r H\}_{r\in\mathbb{Z}}$ is an infinite family of hyperplanes such that $\{g^r H \cap \widetilde{X}_{\widetilde{v}}\}_{r\in\mathbb{Z}}$ is a set of disjoint parallel lines in $\widetilde{X}_{\widetilde{v}}$. As $\widetilde{Y}$ is finite dimensional, there exists an N such that H and $g^N H$ do not intersect. Otherwise $\{g^r H\}_{r\in\mathbb{Z}}$ would be an infinite set of pairwise intersecting hyperplanes, which would imply that there are cubes of arbitrary dimension in $\widetilde{Y}$.

Therefore, as there are only finitely many $G_{\widetilde{v}}$-orbits of hyperplanes intersecting $\widetilde{X}_{\widetilde{v}}$, there exists a finite set of hyperplanes $H_1, \dots, H_d \in \mathcal{H}_{\widetilde{v}}$ and $g_1, \dots, g_d \in G$ such that $\mathcal{H}_{\widetilde{v}} = \{g_1^r H_1, \dots, g_d^r H_d\}_{r\in\mathbb{Z}}$, and each $\{g_i^r H_i\}_{r\in\mathbb{Z}}$ is a disjoint set of hyperplanes

in $\widetilde{Y}$. Therefore $\{g_i^r H_i \cap \widetilde{E}\}_{r\in\mathbb{Z}}$ is a set of disjoint geodesic lines in $\widetilde{X}_{\widetilde{v}}$. Thus, given a 0-cube y, there exists a unique $y_i \in \mathbb{Z}$ such that $y[g_i^{y_i} H_i]$ and $y[g_i^{y_i+1} H_i]$ properly intersect each other. Therefore, construct $\phi : \mathsf{hull}(\widetilde{X}_{\widetilde{v}}) \to \mathbf{R}^d$ by letting $\phi(y) = (y_1, \dots, y_d)$ for each 0-cube y. The map ϕ extends to the 1-skeleton of $\mathsf{hull}(\widetilde{X}_{\widetilde{v}})$ since adjacent 0-cubes lie on the opposite sides of precisely one hyperplane. Therefore ϕ extends to the higher dimensional cubes, and thus $\mathsf{hull}(\widetilde{X}_{\widetilde{v}})$. □

Let $\widetilde{X} \to \widetilde{Y}$ be the lift of the universal cover of an amicable immersion $X \to Y$. Let $\widetilde{X}_{\widetilde{e}}$ be an edge space adjacent to a vertex space $\widetilde{X}_{\widetilde{v}}$. A hyperplane H in $\widetilde{Y}$ intersects $\widetilde{X}_{\widetilde{v}}$ *parallel* to $\widetilde{X}_{\widetilde{e}}$ if $H \cap \widetilde{X}_{\widetilde{v}}$ is a geodesic line parallel to $\widetilde{X}_{\widetilde{e}} \cap \widetilde{X}_{\widetilde{v}}$. Otherwise, if $H \cap \widetilde{X}_{\widetilde{v}}$ is a geodesic line that is not parallel to $\widetilde{X}_{\widetilde{e}} \cap \widetilde{X}_{\widetilde{v}}$, then we say H intersects $\widetilde{X}_{\widetilde{v}}$ *non-parallel* to $\widetilde{X}_{\widetilde{e}}$.

Lemma 5.6. *Let $X \to Y$ be an amicable immersion. Let $\widetilde{e}$ be an edge in $\widetilde{\Gamma}$. Suppose that $H \subseteq \widetilde{Y}$ is a hyperplane intersecting $\widetilde{X}_{-\widetilde{e}}$ non-parallel to $\widetilde{X}_{\widetilde{e}}$, then H intersects $\widetilde{X}_{+\widetilde{e}}$ non-parallel to $\widetilde{X}_{\widetilde{e}}$. Moreover, there is an arc in $H \cap \widetilde{X}_{\widetilde{e}}$ joining $H \cap \widetilde{X}_{-\widetilde{e}}$ to $H \cap \widetilde{X}_{+\widetilde{e}}$.*

Proof. Let $G = \pi_1 X$. The geodesic lines $H \cap \widetilde{X}_{-\widetilde{e}}$ and $\widetilde{X}_{\widetilde{e}} \cap \widetilde{X}_{-\widetilde{e}}$ are non parallel in $\widetilde{X}_{-\widetilde{e}}$, and therefore intersect in a single point $p \in \widetilde{X}_{\widetilde{e}} \cap \widetilde{X}_{-\widetilde{e}}$. As H is two sided in $\widetilde{Y}$ and X the vertex and edge spaces are transverse to H, the intersection of H with X is also locally two sided in $\widetilde{X}$. Therefore, p is contained inside a curve in $H \cap \widetilde{X}_{\widetilde{e}}$. As $\widetilde{X}_{\widetilde{e}}$ is $G_{\widetilde{e}}$-invariant and only finitely many hyperplanes separate any two points in $\widetilde{X}$, we can deduce that p is an endpoint of a compact curve in $H \cap \widetilde{X}_{\widetilde{e}}$ with its other endpoint contained in $\widetilde{X}_{+\widetilde{e}} \cap \widetilde{X}_{\widetilde{e}}$. Thus, H must also intersect $\widetilde{X}_{+\widetilde{e}}$ non-parallel to $\widetilde{X}_{\widetilde{e}}$. □

Lemma 5.7. *Let $X \to Y$ be an amicable immersion, where Y is a finite dimensional, locally finite, nonpositively curved cube complex. If $\pi_1 X \cong \mathbb{Z}^2 *_{\mathbb{Z}} \mathbb{Z}^2$, then for every vertex space $\widetilde{X}_{\widetilde{v}}$ and adjacent edge space $\widetilde{X}_{\widetilde{e}}$ there is a hyperplane H in $\widetilde{Y}$ that intersects $\widetilde{X}_{\widetilde{v}}$ parallel to $\widetilde{X}_{\widetilde{e}}$.*

Proof. Let $G = \pi_1 X$. There are precisely two vertex orbits and one edge orbit in $\widetilde{\Gamma}$. Assume that $\widetilde{Y} = \mathsf{hull}(\widetilde{X})$. Let $\mathcal{H}$ denote the set of all hyperplanes in $\widetilde{Y}$ intersecting $\widetilde{X}$. Let $\mathcal{H}_{\widetilde{v}}$ denote the set of all hyperplanes intersecting $\widetilde{X}_{\widetilde{v}}$.

For each vertex $\widetilde{v}$, there is precisely one $G_{\widetilde{v}}$-orbit of adjacent edges $\widetilde{e}$. Therefore, if $H \in \mathcal{H}_{\widetilde{v}}$ is non-parallel to an adjacent edge space to $\widetilde{X}_{\widetilde{v}}$, it is non-parallel to all adjacent edge spaces, and by Lemma 5.6 it must intersect all adjacent vertex spaces non-parallel to all adjacent edge spaces. Therefore, we deduce that any hyperplane that doesn't intersect every vertex space in $\widetilde{X}$ will either intersect a vertex space

parallel to its adjacent edge spaces, or its intersection will be a line contained in an edge space.

Suppose that there exists a vertex space $\widetilde{X}_{\widetilde{v}_1}$ such that no hyperplane in $\mathcal{H}_{\widetilde{v}_1}$ intersects $\widetilde{X}_{\widetilde{v}_1}$ parallel to the adjacent edge spaces. Let $\hat{\mathcal{H}} = \mathcal{H} - \mathcal{H}_{\widetilde{v}_1}$. Every hyperplane in $\mathcal{H} - \mathcal{H}_{\widetilde{v}_1}$ must intersect each wall in $\mathcal{H}_{\widetilde{v}_1}$ so we deduce that $\widetilde{Y} = \mathsf{hull}(\widetilde{X}_{\widetilde{v}_1}) \times C(\widetilde{Y}, \hat{\mathcal{H}})$ (see [4, Lem 2.5]). Furthermore, $\mathcal{H}_{g\widetilde{v}_1} = \mathcal{H}_{\widetilde{v}_1}$ for all $g \in G$. By Lemma 5.5, $\mathsf{hull}(\widetilde{X}_{\widetilde{v}})$ embeds in $\mathbf{R}^d$. Since $\widetilde{X}_{\widetilde{v}_1}$ is contained inside some subcomplex $\mathsf{hull}(\widetilde{X}_{\widetilde{v}_1}) \times C \subseteq \widetilde{Y}$, where C is the 0-cube determined by orienting all hyperplanes towards $\widetilde{X}_{\widetilde{v}_1}$, we can conclude that only finitely many hyperplanes in $\hat{\mathcal{H}}$ intersect the r-neighborhood of $\widetilde{X}_{\widetilde{v}_1}$.

Let $\widetilde{X}_{\widetilde{u}}$ be an vertex space adjacent to $\widetilde{X}_{\widetilde{v}_1}$, and let $\widetilde{X}_{\widetilde{e}_1}$ be the edge space connecting them. Let $\widetilde{X}_{\widetilde{v}_2}$ be another vertex space adjacent to $\widetilde{X}_{\widetilde{u}}$ and let $\widetilde{X}_{\widetilde{e}_2}$ be the edge space connecting them. Note that $\widetilde{X}_{\widetilde{v}_1}$ and $\widetilde{X}_{\widetilde{v}_2}$ are in the same $G_{\widetilde{u}}$-orbit. The geodesic lines $\widetilde{X}_{\widetilde{u}} \cap \widetilde{X}_{\widetilde{e}_1}$ and $\widetilde{X}_{\widetilde{u}} \cap \widetilde{X}_{\widetilde{e}_2}$ are parallel in $\widetilde{X}_{\widetilde{u}}$. Let $D \subseteq \widetilde{X}_{\widetilde{u}}$ be the subspace isometric to $\mathbb{R} \times [a, b]$ bounded by these parallel lines. Let $U = D \cup \widetilde{X}_{\widetilde{e}_1} \cup \widetilde{X}_{\widetilde{e}_2}$. Finitely many hyperplanes in $\hat{\mathcal{H}}$ intersect D.

Let g_1 be an isometry in $G_{\widetilde{v}_1}$ that stabilizes an axis in $\widetilde{X}_{\widetilde{v}_1}$ that is non-parallel to the geodesic $\widetilde{X}_{\widetilde{v}_1} \cap \widetilde{X}_{\widetilde{e}_1}$. Similarly, let g_2 be an isometry in $G_{\widetilde{v}_2}$ that stabilizes an axis in $\widetilde{X}_{\widetilde{v}_2}$ that is non-parallel to the geodesic $\widetilde{X}_{\widetilde{v}_2} \cap \widetilde{X}_{\widetilde{e}_2}$. Note that $F = \langle g_1, g_2 \rangle$ is a free group on two generators. Let r be such that U is contained in the r-neighbourhood of $\widetilde{X}_{\widetilde{v}_1}$. As there are only finitely many hyperplanes in $\hat{\mathcal{H}}$ intersecting $\mathcal{N}_r(\widetilde{X}_{\widetilde{v}_1})$, there must exist an n such that g_1^n stabilizes those walls. Similarly, since $\widetilde{X}_{\widetilde{v}_2}$ is a translate of $\widetilde{X}_{\widetilde{v}_1}$, we can deduce that there are only finitely many hyperplanes in $\hat{\mathcal{H}}$ intersecting $\mathcal{N}_r(\widetilde{X}_{\widetilde{v}_2})$, and there must exist an m such that g_2^m stabilizes those walls. Let $F' = \langle g_1^n, g_2^m \rangle$. As U lies in both $\mathcal{N}_r(\widetilde{X}_{\widetilde{v}_1})$ and $\mathcal{N}_r(\widetilde{X}_{\widetilde{v}_2})$ we can deduce that the hyperplanes in $\hat{\mathcal{H}}$ that intersect the F'-translates of U are precisely the hyperplanes intersecting U. Let $\widetilde{Z} = F'\widetilde{X}_{\widetilde{v}_1} \cup F'\widetilde{X}_{\widetilde{v}_2} \cup F'U$. Then $\mathsf{hull}(\widetilde{Z}) = \mathsf{hull}(\widetilde{X}_{\widetilde{v}_1}) \times K \subseteq \widetilde{Y}$, where K is a compact cube complex. Then F acts freely on $\mathsf{hull}(\widetilde{Z})$, which is a contradiction since number of 0-cubes intersecting the r-neighborhood of a 0-cube in $\mathsf{hull}(\widetilde{X}_{\widetilde{v}_1}) \times K$ grows polynomially with r, and therefore cannot permit a free F-action. □

Lemma 5.7 is a special case of the following more general statement:

Corollary 5.8. *Let $X \to Y$ be an amicable immersion, where Y is a finite dimensional, locally finite, nonpositively curved cube complex. Then for every vertex space $\widetilde{X}_{\widetilde{v}}$ and adjacent edge space $\widetilde{X}_{\widetilde{e}}$ there is a hyperplane H in $\widetilde{Y}$ that intersects $\widetilde{X}_{\widetilde{v}}$ parallel to $\widetilde{X}_{\widetilde{e}}$.*

Proof. For every edge $\widetilde{e}$ in $\widetilde{\Gamma}$ there is a subgroup $G' = \langle G_{-\widetilde{e}}, G_{+\widetilde{e}} \rangle \leq G$ such that $G' \cong \mathbb{Z}^2 *_{\mathbb{Z}} \mathbb{Z}^2$. Let $Y' = G' \backslash \widetilde{Y}$. Then there is an amicable immersion $X' \to Y'$ such that $\widetilde{X}'_{\widetilde{e}} = \widetilde{X}$ and $\widetilde{X}'_{\pm\widetilde{e}} = \widetilde{X}_{\pm\widetilde{e}}$. Therefore, by Lemma 5.7, there is a hyperplane intersecting $\widetilde{X}_{-\widetilde{e}}$ parallel to $\widetilde{X}_{\widetilde{e}}$, and similarly for $\widetilde{X}_{+\widetilde{e}}$. □

The following proposition is a strengthening of one direction of Theorem 1.1 in [11]. Let $f_1 : A \to C$ and $f_2 : B \to C$ be maps between topological spaces A, B, C. The *fiber product* $A \otimes_C B = \{(a, b) \in A \times B \mid f_1(a) = f_2(b)\}$. Note that there are natural projections $p_1 : A \otimes_C B \to A$ and $p_2 : A \otimes_C B \to B$.

Proposition 5.9. *Let G be a tubular group acting freely on a CAT(0) cube complex $\widetilde{Y}$. Then there is a tubular space X with a finite set of immersed walls such that the associated wallspace $(\widetilde{X}, \mathcal{W})$ has the following properties:*

1. *G acts freely on $C(\widetilde{X}, \mathcal{W})$.*
2. *$C(\widetilde{X}, \mathcal{W})$ is finite dimensional if $\widetilde{Y}$ is finite dimensional.*
3. *$C(\widetilde{X}, \mathcal{W})$ is finite dimensional and locally finite if $\widetilde{Y}$ is locally finite.*

Proof. Let $Y = G \backslash \widetilde{Y}$. Let $X \to Y$ be an amicable immersion. Assume that $\widetilde{Y} = \mathsf{hull}(\widetilde{X})$, so every immersed hyperplane in Y intersects X. Therefore, as X is compact, there are finitely many immersed hyperplanes $h_1, \dots h_m$ in Y.

Let $h_i \to Y$ be an immersed hyperplane in Y. We obtain horizontal immersed walls in X by considering the components of the fiber product $X \otimes_Y h_i$ of $X \to Y$ and $h \to Y$. Each component Λ has a natural map into X. The components of $X \otimes_Y h$ that have image in X contained in an edge space are ignored. Let Λ^p be a component of $X \otimes_Y h$ whose image in X intersects a vertex space $X_v \subseteq X$. We will show that after a minor adjustment to Λ^p, we obtain a horizontal immersed wall and by considering all such components we obtain a set of horizontal walls in X obtained from an equitable set.

Using the map $\Lambda^p \to X$ we can decompose Λ^p into the components of the preimages of vertex space and edge spaces. As the intersection of each hyperplane $H \subseteq \widetilde{Y}$ with each vertex space $\widetilde{X}_{\widetilde{v}}$ is either empty or a geodesic line, the intersection of each h_i with X_v is a set of geodesic curves, so Λ^p restricted to the preimage of X_v is a set of geodesic curves. By Lemma 5.6 each hyperplane $H \subseteq \widetilde{Y}$ that intersects a vertex space $X_{\widetilde{v}}$ non-parallel to an adjacent edge space $\widetilde{X}_{\widetilde{e}}$ will intersect $\widetilde{X}_{\widetilde{e}}$ as an arc with endpoints in $\widetilde{X}_{-\widetilde{e}}$ and $\widetilde{X}_{+\widetilde{e}}$. Thus the components of the intersection $X_e \cap h_i$ that intersect X_{-e} or X_{+e}, are arcs with endpoints in both X_{-e} and X_{+e}. Therefore, Λ^p decomposes into circles that map as local geodesics into vertex spaces, and arcs that map into edges spaces X_e with an endpoint in each X_{-e} and X_{+e}.

Let $\{\Lambda_1^p, \dots, \Lambda_k^p\}$ be the set of all such components of $X \otimes_Y h_i$ that intersect vertex spaces. Let S_v^p be the set curves that map the circles in $\{\Lambda_i^p\}_{i=1}^k$ to the vertex space X_v. The elements of S_v^p and the attaching maps $\varphi_e^{\pm}$ of the edge spaces in X are locally geodesic curves, and $\#[\varphi_e^+, S_{e^+}] = \#[\varphi_e^-, S_{e^-}]$ since both sides are equal to the number of arcs in the walls $\{\Lambda_i^p\}_{i=1}^k$ that map into X_e. As G acts freely on $\widetilde{Y}$, there must be hyperplanes intersecting each vertex space $\widetilde{X}_{\widetilde{e}}$ as geodesics in at least two parallelism classes. This implies that S_v^p contains curves generating at least two non-commensurable cyclic subgroups of G_v, and therefore S_v^p generates a finite index subgroup of G_v.

S_v^p is almost an equitable set: the images of the curves in S_v^p may not be disjoint. Suppose that $\alpha_1, \dots \alpha_m \in S_v$ be a maximal set of curves that have identical image in X_v. Let $\mathcal{N}_\epsilon(Q)$ denote the ϵ-neighborhood of a subset Q of either Y or $\widetilde{Y}$ with respect to the CAT(0) metric. Let $\epsilon \in (0, \frac{1}{3})$ be such that the neighbourhood $\mathcal{N}_\epsilon(\alpha_1) \subseteq Y$ only contains the images of $\alpha_1, \dots, \alpha_m$ and the arcs connected to them. There is a homotopy of $\bigsqcup_{i=1}^k \Lambda_i^p \to X$ that is the identity outside of $\bigsqcup_{i=1}^k \Lambda_i^p \cap N_\epsilon(\alpha_1)$ such that $\alpha_1, \dots, \alpha_m$ are homotoped to a disjoint set of geodesic curves in $X_v \cap N_\epsilon(\alpha_1)$ transverse or disjoint from all the other curves in S_v^p. By choosing ϵ small enough we can perform such a homotopy $\Phi : \bigsqcup_{i=1}^k \Lambda_i^p \times [0,1] \to X$ such that all sets of overlapping curves in $\{S_v^p\}_{v \in V\Gamma}$ become disjoint and such that Φ is the identity map outside of the ϵ-neighborhood of the overlapping curves. The restriction of Φ to $\Lambda_i^p \times \{1\} \to X$ is an immersed wall that we will denote by Λ_i. Thus, the immersed walls $\{\Lambda_i\}_{i=1}^k$ obtained from an equitable set $\{S_v\}_{v \in V\Gamma}$. We refer to $\{\Lambda_i^p\}_{i=1}^k$ as the *immersed proto-walls* and the lifts $\widetilde{\Lambda}_i^p \to \widetilde{X}$ as the *proto-walls*. Note that proto-walls have regular and non-regular intersections in the same way that walls do.

Let $(\widetilde{X}, \mathcal{W})$ be the wallspace obtained from the immersed walls $\{\Lambda_i\}_{i=1}^k$ and adding a single vertical wall for each edge space. Each wall $\widetilde{\Lambda} \to \widetilde{X}$ covers an immersed wall $\Lambda \to X$. There exists a homotopy of $\Lambda \to X$ to the corresponding immersed proto-wall $\Lambda^p \to X$. This homotopy lifts to a homotopy from the immersed wall $\widetilde{\Lambda} \to \widetilde{X}$ to a unique proto-wall $\widetilde{\Lambda}^p \to \widetilde{X}$. Note that each wall is contained in the ϵ-neighborhood of its corresponding proto-wall. Each proto-wall corresponds to the intersection of a unique hyperplane in $\widetilde{Y}$ with the image of $\widetilde{X}$ in $\widetilde{Y}$. Therefore, each wall in $\mathcal{W}$ corresponds to a unique hyperplane in $\widetilde{Y}$.

Let $\widetilde{\Lambda}$ be a wall in $\mathcal{W}$, and let $\widetilde{\Lambda}^p$ be the corresponding proto-wall. Note that $\widetilde{\Lambda} \cap \widetilde{X}_{\widetilde{v}}$ and $\widetilde{\Lambda}^p \cap \widetilde{X}_{\widetilde{v}}$ are either parallel geodesic lines, or both empty intersections. Therefore, if $\widetilde{\Lambda}_1, \widetilde{\Lambda}_2 \in \mathcal{W}$ are a pair of regularly intersecting walls, then they correspond to a pair of regularly intersecting proto-walls, which correspond to a pair of intersecting hyperplanes in $\widetilde{Y}$.

If a pair of proto-walls $\widetilde{\Lambda}_1^p$ and $\widetilde{\Lambda}_2^p$ are disjoint, then the corresponding walls in $\widetilde{\Lambda}_1$ and $\widetilde{\Lambda}_2$ in $\mathcal{W}$ are also disjoint. Moreover, since $\widetilde{\Lambda}$ is contained in the ϵ-neighborhood of $\widetilde{\Lambda}^p$, a halfspace of $\widetilde{\Lambda}$ determines a halfspace of $\widetilde{\Lambda}^p$ and therefore a halfspace of the hyperplane H corresponding to $\widetilde{\Lambda}^p$.

To prove (2), suppose that $C(\widetilde{X}, \mathcal{W})$ were infinite dimensional, then by Proposition 2.2, there would exists an infinite set of pairwise regularly intersecting walls in $\mathcal{W}$, which implies there is an infinite set of pairwise regularly intersecting proto-walls. Therefore, there is an infinite set of pairwise intersecting hyperplanes in $\widetilde{Y}$. This would imply that $\widetilde{Y}$ is an infinite dimensional CAT(0) cube complex. Therefore, if $\widetilde{Y}$ is finite dimensional, then so is $C(\widetilde{X}, \mathcal{W})$.

To prove (3) we first prove the following:

Claim. *If $\widetilde{Y}$ is locally finite, then $C(\widetilde{X}, \mathcal{W})$ is finite dimensional.*

Proof. Suppose that $\widetilde{Y}$ is locally finite. If $C(\widetilde{X}, \mathcal{W})$ is infinite dimensional, then by Lemma 2.2 it contains an infinite cube containing a canonical 0-cube z. Let $\widetilde{\Lambda}_1, \dots, \widetilde{\Lambda}_n, \dots$ be the set of infinite pairwise crossing walls corresponding to the infinite cube. Let $\widetilde{\Lambda}_1^p, \dots, \widetilde{\Lambda}_n^p, \dots$ be the corresponding set of infinite pairwise crossing proto-walls, and let $H_1, \dots, H_n, \dots$ be the corresponding infinite family of pairwise crossing hyperplanes.

Suppose that Q is a subcomplex in $\widetilde{Y}$. Let $U(Q)$ denote the *cubical neighborhood* of Q, which is the union of all cubes in $U(Q)$ that intersect Q. As $\widetilde{Y}$ is locally finite, if Q is compact, then $U(Q)$ is also compact. By [7, Lem 13.15], if Q is convex, then so is $U(Q)$. Let $U^n(Q)$ denote the cubical neighborhood of $U^{n-1}(Q)$.

Let $x \in \widetilde{X}$ be a point determining the canonical 0-cube z in $C(\widetilde{X}, \mathcal{W})$. Let x be contained in a cube C in $\widetilde{Y}$. As C is compact and convex, $U^n(C)$ is also compact and convex, and therefore can only be intersected by finitely many H_i. There exists an H_i such that H_i intersects $U^{N+2}(C)$, but not $U^{N+1}(C)$ for some $N > 1$. Since H_i does not intersect $U^N(C)$ nor $U^{N+1}(C)$, there must exist a hyperplane H intersecting $U^{N+1}(C)$ that separates $U^N(C)$ from H_i. Note that $\mathsf{d}_{\widetilde{Y}}(x, H) \geq N$ and $\mathsf{d}_{\widetilde{Y}}(H, H_i) \geq 1$. Let $\widetilde{\Lambda}^p$ be the proto-wall corresponding to H and $\widetilde{\Lambda}$ be the corresponding wall. As $\widetilde{\Lambda}^p$ separates x from $\widetilde{\Lambda}_i^p$, we can conclude $\widetilde{\Lambda}$ separates x from $\widetilde{\Lambda}_i$ since the $\widetilde{\Lambda}$ and $\widetilde{\Lambda}_i$ are respectively contained in the ϵ-neighborhoods of $\widetilde{\Lambda}^p$ and $\widetilde{\Lambda}_i^p$. This contradicts the fact that z is incident to a 1-cube dual to hyperplane corresponding to $\widetilde{\Lambda}_i$. □

As $C(\widetilde{X}, \mathcal{W})$ is finite dimensional, we can apply Corollary 5.8 to each edge group in G to deduce that $\{\Lambda_i\}_{i=1}^k$ are fortified. Therefore, by Proposition 4.7 we deduce that $C(\widetilde{X}, \mathcal{W})$ is locally finite. □

We can now prove the main theorem of this paper.

Theorem 5.10. *A tubular group G acts freely on a locally finite CAT(0) cube complex if and only if G is virtually special.*

Proof. Suppose that G is virtually special. Then G embeds as the subgroup of a finitely generated right angled Artin group, and therefore acts freely on the universal cover of the corresponding Salvetti complex, which is necessarily locally finite.

Conversely, suppose that G acts freely on a locally finite CAT(0) cube complex. Let X be a tubular space such that $G = \pi_1 X$. By Proposition 5.9 there exists a finite set of immersed walls such that the dual of the associated wallspace $C(\widetilde{X}, \mathcal{W})$ is finite dimensional and locally finite. By Lemma 3.1 we can assume that the immersed walls are also primitive. Therefore, by Proposition 4.8, G is virtually special. □

6. Virtual Cubical Dimension

Lemma 6.1. *Let X be a tubular space and $G = \pi_1 X$. Suppose there exists an equitable set that produces primitive, non-dilated immersed walls in X. There exists a finite index subgroup $G' \leqslant G$ such that for each vertex group G'_v of the induced splitting of G' the natural map $\mathsf{H}_1(G'_v) \to \mathsf{H}_1(G')$ is an injection as a summand.*

Proof. Note that $\mathsf{H}_1(G)$, and its finite index subgroups, is a summand of two factors. The first factor $\mathsf{H}_1^{\mathsf{v}}(G)$ is generated by the image of the vertex groups of G, and the second factor $\mathsf{H}_1^{\mathsf{v}}(G)$ is generated by the stable letters in the graph of groups presentation.

Let $\widetilde{\Gamma}$ be the Bass-Serre tree. By Proposition 4.6, since there are immersed walls that are primitive and non-dilated, G acts freely on $\mathbf{R}^d \times \widetilde{\Gamma}$ such that $G_{\widetilde{v}}$ fixes the vertex $\widetilde{v}$ in $\widetilde{\Gamma}$. Therefore G is a subgroup of $\mathrm{Aut}(\mathbf{R}^d \times \widetilde{\Gamma}) \cong \mathbb{Z}^d \rtimes \mathrm{Aut}([-1,1]^d) \times \mathrm{Aut}(\widetilde{\Gamma})$. There is a finite quotient $\mathbb{Z}^d \rtimes \mathrm{Aut}([-1,1]^d) \times \mathrm{Aut}(\widetilde{\Gamma}) \to \mathrm{Aut}([-1,1]^d)$, so let $G'' \leqslant G$ be the finite index subgroup contained in the kernel. Note that G'' embeds in $\mathbb{Z}^d \times \mathrm{Aut}(\widetilde{\Gamma})$. Let $p : G'' \to \mathbb{Z}^d$ be the projection onto the first factor. Each vertex group survives in the image of p, and therefore we have embedding $\mathsf{H}_1(G''_v) \hookrightarrow \mathsf{H}_1^{\mathsf{v}}(G'') \hookrightarrow \mathbb{Z}^d$.

For each vertex v in $\Gamma'' = \widetilde{\Gamma}/G''$ there is a finite index subgroup $A_v \leqslant \mathbb{Z}^d$ such that $p(G''_v) \cap A_v$ is a summand of A_v. Let $A = \cap_v A_v$ and $G' = p^{-1}(A)$. Each vertex group in G' will be a factor in A. As A is free abelian, the map $G' \to A$ will factor through the $\mathsf{H}_1^{\mathsf{v}}(G')$, so we can deduce that each vertex group survives as a retract in $\mathsf{H}_1(G')$. Therefore, each vertex group in G' survives as a summand in the first homology. □

Theorem 6.2. *Let G be a tubular group acting freely on a finite dimensional CAT(0) cube complex. Then there is a finite index subgroup $G' \leqslant G$ such that G' acts freely on a 3-dimensional CAT(0) cube complex.*

Proof. Let X be a tubular space such that $G = \pi_1 X$. By Proposition 5.9 there exists immersed walls in X that are non-dilated, and by Lemma 3.1 we can assume that they are also primitive. Let $G' \leqslant G$ be the finite index subgroup given by Lemma 6.1, and let X' be the corresponding covering space. Let $\mathsf{H}_1^{\mathsf{v}}(G') \cong \mathbb{Z}^d$ be the summand in the first homology generated by the vertex groups. By Lemma 6.1, there is an inclusion $\iota_v : G'_v \to \mathsf{H}_1^{\mathsf{v}}(G')$, and a projection map $p_v : \mathsf{H}_1^{\mathsf{v}}(G') \to G'_v$.

Suppose that $d = 2$. Choose any pair of elements $a, b \in \mathsf{H}_1^{\mathsf{v}}(G')$ that generate $\mathsf{H}_1^{\mathsf{v}}(G')$. We claim that $S = \{S_v = \{p_v(a), p_v(b)\} \mid v \in V(\Gamma)\}$ is an equitable set. By construction, each S_v generates G'_v. The edge group $G'_e = \langle g_e \rangle$ is adjacent to G'_{-e} and G'_{+e} and the respective inclusions are $(\varphi_e^-)_*$ and $(\varphi_e^+)_*$. There is an isomorphism $A = p_{+e} \circ \iota_{-e} \in SL(2, \mathbb{Z})$, that maps $(\varphi_e^-)_*(g_e)$ in G'_u to $(\varphi_e^+)_*(g_e)$ in G'_v. Therefore,

$$\#[p_u(a), (\varphi_e^-)_*(g_e)] = \#[Ap_u(a), A(\varphi_e^-)_*(g_e)] = \#[p_v(a), (\varphi_e^+)_*(g_e)],$$

A congruent set of equalities exist for b. These equalities also imply that the choice of arcs for the equitable sets can be chosen so that they only join circles that are the image of the same elements of $\mathsf{H}_1^{\mathsf{v}}(G')$. Therefore a set of embedded immersed walls is obtained with precisely two immersed walls intersecting each vertex space: such a set of horizontal immersed walls along with a vertical wall for each edge will give a three dimensional dual cube complex.

If $d > 2$, then there exist vertex groups G_u and G_v such that they embed into $\mathsf{H}_1^{\mathsf{v}}(G')$ as distinct summands. Let $G_u = \langle g_u, g'_u \rangle$ and $G_v = \langle g_v, g'_v \rangle$. We can assume, since they are distinct summands that g_u is disjoint from the image of G_v in $\mathsf{H}_1^{\mathsf{v}}(G')$, and that g_v is disjoint from the image of G_u in $\mathsf{H}_1^{\mathsf{v}}(G')$. By attaching an edge space in X' connecting X_u and X_v, that have attaching maps representing g_u and g_v respectively we obtain a new graph of spaces $\hat{X}'$. The resulting tubular group $\hat{G}' = \pi_1 \hat{X}'$ has $\mathsf{H}_1^{\mathsf{v}}(\hat{G}') \cong \mathbb{Z}^{d-1}$, so by induction we can obtain the specified graph of spaces $\hat{X}'$. Given a set of immersed walls for $\hat{X}'$ with a dual cube complex of dimension at most 3 we obtain immersed walls for X' by deleting the arcs that map into the the edge space that was added to construct $\hat{X}'$. The immersed walls obtained still give a dual cube complex with dimension at most 3. □

References

[1] N. Brady and M. R. Bridson. There is only one gap in the isoperimetric spectrum. *Geom. Funct. Anal.*, 10(5):1053–1070, 2000.

[2] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[3] J.O. Button. Tubular free by cyclic groups and the strongest tits alternative.

[4] Pierre-Emmanuel Caprace and Michah Sageev. Rank rigidity for CAT(0) cube complexes. *Geom. Funct. Anal.*, 21(4):851–891, 2011.

[5] Christopher H. Cashen. Quasi-isometries between tubular groups. *Groups Geom. Dyn.*, 4(3):473–516, 2010.

[6] Frédéric Haglund and Frédéric Paulin. Simplicité de groupes d'automorphismes d'espaces à courbure négative. In *The Epstein birthday schrift*, volume 1 of *Geom. Topol. Monogr.*, pages 181–248 (electronic). Geom. Topol. Publ., Coventry, 1998.

[7] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *Geom. Funct. Anal.*, 17(5):1551–1620, 2008.

[8] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.

[9] Daniel T. Wise. Research announcement: the structure of groups with a quasiconvex hierarchy. *Electron. Res. Announc. Math. Sci.*, 16:44–55, 2009.

[10] Daniel T. Wise. *From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry*, volume 117 of *CBMS Regional Conference Series in Mathematics*. Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 2012.

[11] Daniel T. Wise. Cubular tubular groups. *Trans. Amer. Math. Soc.*, 366(10):5503–5521, 2014.

[12] D.T. Wise and C. Hruska. Finiteness properties of cubulated groups. Submitted for publication.

[13] D.J. Woodhouse. Classifying finite dimensional cubulations of tubular groups. Submitted for publication.

[14] D.J. Woodhouse. A generalized axis theorem for cube complexes. Submitted for publication, 2016.

E-mail address: daniel.woodhouse@mail.mcgill.ca